

Project UNITY: Cross Domain Visualization Collaboration

Jason Moore

Air Force Research Laboratory, Information Directorate

Dr. Paul Havig

Air Force Research Laboratory, 711th Human Performance Wing

ABSTRACT

Distributed collaboration research and applications have created the capability for documents, images, whiteboards and computer displays to be shared by one of more collaborators. Commercial offerings from companies like Google, Adobe, Microsoft, IBM, and others offer a simplistic subset of access protocols and sharing methodologies. While these methods work effectively for their purposes, the Department of Defense and their partners have limited capability to use them due to policy and security issues, which are especially stressed during coalition operations. This paper describes visualization collaboration issues and solutions for the forthcoming international experiment, project UNITY. The major difference between traditional Relaxed What You See Is What I See systems available today that provide a separate private workspace is that the UNITY project situates private content in-situ with shared content while ensuring that the operators understand the difference.

1. INTRODUCTION

UNITY is an International Cooperative Research and Development (ICR&D) project between the United States and Great Britain under the Research and Development Projects (RDP) Memorandum of Agreement (MOA). UNITY's objectives are to develop and evaluate the operational concepts and requirements for undertaking combined operations pursuant to the interests of mission partners. The partners agreed to develop, experiment, and demonstrate transitionable emergent technologies, capabilities, or concepts that facilitate the sharing of information and products. The Air Force Research Laboratory (AFRL) created an end-to-end system by combining and enhancing a suite of software tools to meet the UNITY objectives. This paper specifically focuses on the visualization and interaction opportunities and technical challenges for collaboration when information cannot be shared uniformly among all interested parties.

1.1 Requirements

Collaboration between coalition partners is essential for accurate and timely decision making in the ever increasing intensity and tempo of global security operations. There are two types of requirements that frame the potential set of solutions for cross domain visualization collaboration: those that are imposed by policy and those that are imposed by users' requirements for flexibly sharing information and sustaining individual and shared Situation Awareness (SA).

1.1.1 Policy Requirements

The Department of Defense (DoD) has strict rules and policies that dictate how data is classified and how networks at differing classifications can be connected. These requirements are primarily specified by:

- *Executive Order 13526 (EO13526)*[17], "...prescribes a uniform system for classifying, safeguarding, and declassifying national security information", across national security disciplines, networks, services, and data.
- *Defense Information System Network (DISN): Policy, Responsibilities and Processes (DISN-PRP)*[12], "...provides guidance on cross domain connections between security domains (i.e., internal to DOD and foreign) and cross functional connections with non-DOD organizations (e.g., non-DOD US Government agencies and contractors) (Enclosure C)".

1.1.2 Situation Awareness Requirements

The user needs two types of spaces: a private space to maintain their individual SA and a public space for shared SA. Furthermore, data available in both spaces needs to be continuously updated and available at machine speed. It is critical that the user be able to, at all times, distinguish between data that they are actively sharing and data that is not shared due to policy or desire. Additionally, data from each participant must be composable into a single visualization or set of coordinated visualizations.

2. ASSUMPTION

The assumption is that everyone in the collaboration is a cooperative client, meaning that collaborators¹ are there to achieve a common goal while adhering to legal and policy requirements. This experiment does not address the intentional misuse of data, like hidden steganographic messages or intentionally mismarked classified information. Other technology (discussed below) is in place to prevent the accidental leakage of classified information whereas the UNITY effort is about the engineering solutions of working within policy across accredited systems.

3. METHODOLOGY

The UNITY project is not developing new cross domain solutions or changing Air Force policy, but rather: *a*) using an existing cross domain solution, the AFRL-developed Information Support Server Environment (ISSE) Guard [8], *b*) modifying the data distribution system Phoenix'[2] to support the unique requirements of communication and in-order delivery over the ISSE Guard, and *c*) enhancing the visualization technology provided by the AFRL Information Directorate-developed Air, Space, and Cyber User-Defined Operational Picture (ASC-UDOP) [1], [3], [15], [16]. This infrastructure creates an environment to support human collaboration over approved machine-to-machine communication methods while preserving the following list of collaboration and interaction requirements.

Users must be able to:

- access the same or similar information as other users
- selectively share information with any set of other users
- locally view both shared and non-shared information within a single canvas while preserving that distinction visually
- share their mouse pointer and visualize shared pointer(s) regardless of rendering differences, such as screen size or aspect ratio
- share visual mappings with other users
- share and create 3D Views and Tabular Views with collaborators
- chat with other users
- use the system in a master-slave mode or a peer-to-peer mode

4. SYSTEM SETUP

4.1 Hardware Setup

The setup currently includes two locations: the research and development site is at the AFRL Information Directorate in Rome NY whereas the space operator location is at the Space Operations Coordination Center (UK-SPOCC) in High Wycombe, UK. Identical AFRL-developed ErgoWorkstations (see Figure 2) were installed in both locations. The AFRL ErgoWorkstation is made up of a high performance Windows-based PC with three displays: two 30" Dell Cinema displays on the ends of a 56" Barco LC-5621 4k display, producing a display system of nearly 16.5 million pixels. The Operational View (OV-1) for the entire system can be seen in Figure 1.

The intent of using identical hardware is to minimize complexity, to simplify debugging, and to provide an opportunity to experiment with a master-slave mode. This additionally readily focuses the research upon interaction methods and

¹The term collaborator has a negative connotation in numerous cultures, and since this is for an international audience, the word collaborant is used.

communication bandwidth limitations rather than on the asymmetries that are inherent in trying to share content between vastly different devices. While this setup consists of only two systems, the software has no limitations and offers the same opportunities independent of the number of participants and disparities in classification.

There are actually two different Virtual Private Networks (VPNs) connecting the two systems: one network that has the ISSE Guard inline and another network connection for management that is used to debug issues, to setup tests, and to compare the UNITY approach against pixel-scraping software (described in more detail in section 5.1) and against the integrated Microsoft Windows Remote Desktop capability.




Figure 1: UNITY-OV1 Concept



Figure 2: Ergonomic Workstation as installed at AFRL, Information Directorate

4.2 Data Setup

In order to test the system, unclassified data was used and marked with higher classifications. **Any classification markings in this document are strictly for illustrative purposes only and all data depicted is notional and UNCLASSIFIED.**

4.3 Software Setup

Both computers are configured with *a)* Windows 7 with Remote Desktop enabled, *b)* a simple chat client using the management network, *c)* Virtual Network Computing (VNC) which also uses the management network, *d)* Java 8, and

e) the ASC-UDOP application.

4.3.1 ASC-UDOP Introduction

The Air, Space and Cyber User-Defined Operational Picture (ASC-UDOP) is an AFRL Information Directorate-developed visualization application that allows users to graphically configure and manipulate data streams from a variety of sources pertaining to the air, space and cyber domains. It applies composable visualization techniques to provide an environment appropriate for dealing with large and varied types of data, ranging from synthetic to live, and doing so with unprecedented configurability for DoD applications.

The ASC-UDOP application provides the user with display containers, transformations, and renderers. A display container is a blank canvas or view that allows the user to add visualizations in order to display the data in the system with user-selected attributes. One of the canonical use cases is to view the live air picture that AFRL has access to via the Federal Aviation Administration (FAA) (see Figure 3). The data source is mapped to the model renderer fields which are labeled: latitude, longitude and altitude. By doing this, the model renderer knows where to place the plane model in the 3D space. Additionally, the user could modify the values of data using transformations. The user could also modify the color of the icons with a transfer-function that maps the altitude to specific color bands, helping to differentiate important air space levels.




Figure 3: Using the ASC-UDOP to display the live air picture over the United States and label the aircraft with the Mode3ID and the radar that is tracking them currently

The pipeline for using the ASC-UDOP is as follows: 1) load data, 2) drag data sources onto the composition window, 3) drag transformations into the composition window, 4) connect data sources to transformations, 5) configure the transformations, 6) drag a display container onto the composition window, 7) drag renderers into the display container, 8) connect the transformations to the renderers, 9) and configure the renderers to display the data as desired. Alternatively, the user can load pre-saved visualizations that have already been created.

For the purpose of the UNITY effort, only the tabular display container (Tabular View) and the 3D display container (3D View) were extended to support the cross domain collaboration. Since chat is already by nature a shared resource, the team incorporated two already approved cross domain chat clients, TransVerse (by APAN) and JChat (an open source tool). The web page viewer and 2D charting offer interesting problems, but are beyond the scope of this effort and represent a smaller subset of issues. The challenges and changes to the 3D and Tabular Views are described in Section 6.

In order to scale up to large and dynamic datasets, a tuple-based data marshaling system was adopted that is loosely based on a project called Rivet[4]. Each row of data is independently sent through the user-created transformation network until it finally reaches the visualization component or components. The marshaling system has an in-order guarantee that is necessary for the ASC-UDOP application to maintain interactive frame rates when presented with large datasets. This means that tuple events (remove, add, or change) are guaranteed to happen in order, which allows for a much smaller number of events to propagate through the system. This in-order requirement posed a problem for the Phoenix' system initially and the necessary changes are described in Section 6.

5. COLLABORATION

5.1 How others address collaboration

A review of team collaboration tools used in the military and government[21] in 2007 documented both Commercial Off The Shelf (COTS) and Government Off The Shelf (GOTS) solutions. Few of those tools even mentioned security as an issue. For those that did, they limited the solutions to: encrypted communication, enterprise authentication, or eliminating file transfer to reduce the spread of malicious code; they addressed security by limiting functionality. This approach will continue to harm the DoD and produce significant barriers to the timely execution of mission objectives.

There are two major ways in which collaborative features are enabled today: *a)* using an additional application that provides pixel-scraping of a participant’s desktop (where the server software copies the contents of the screen and transmits it to the client) or *b)* integrated as an application workspace feature. The sheer number of open-source and commercial offerings is staggering. Pixel-scraping options include relatively simple desktop cloning software, such as Virtual Network Computing (VNC), in contrast with enterprise offerings like Adobe Connect that offer instant cross firewall collaboration, rich multimedia presentations, user authentication, secure communication channels, thin-client support, and more. Shared workspace methods come in the form of a whiteboard, spreadsheet, or notepad in offerings like Etherpad, Google Docs, or Google Sheets.

One thing pixel-scraping and shared workspace solutions have in common is the concept of shared and private spaces. In pixel-scraping tools, private spaces are defined by segmenting part of the screen or by sharing only a single window. However, a major limitation in these tools is the all-or-nothing binary segmentation of what is being shared versus not shared. There is a distinct lack of capability to maintain one’s unique SA in context with the group’s SA.

A recent publication by Greenberg et al [13] starts to recognize the challenges and desires to have portions of the content shared with different participants, “both sides of the display must be able to present different content, albeit selectively.” The transparent display proposed by Li et al[13] discusses the needs to have a private area of work. While they recognize the value of private spaces; the collocation of shared and private spaces is still not addressed.

The goal of collaboration is workspace awareness as explained by Gutwin and Greenburg in [11]. Through providing a shared context, each user can answer the *who*, *what*, and *where* questions to support the following activities: *management of coupling*, *simplification of communication*, *coordination of action*, *anticipation*, and *assistance*. Table 1 presents the activities where Workspace Awareness (WA) is used and Table 4 presents the concepts and terms for WA that are relevant to the UNITY program.

Table 1: “Summary of the activities in which workspace awareness is used” from [11]

Activity	Benefit of workspace awareness
Management of coupling	Assist people in noticing and managing transitions between individual and shared work.
Simplification of communication	Allows people to use the workspace and artifacts as conversational props, including mechanisms of deixis [words where context is necessary as in “yesterday” and “here”], demonstrations, and visual evidence.
Coordination of action	Assists people in planning and executing low-level workspace actions to mesh seamlessly with others.
Anticipation	Allows people to predict others’ actions and activity at several time scales.
Assistance	Assists people in understanding the context where help is to be provided.

Gutwin and Greenburg also describe two types of placement: *situated* and *separate* and two types of presentation: *literal* and *symbolic*. *Situated* placement deals with information that is displayed within the workspace where *separate* is information that is not depicted within the workspace. *Literal* displays show information in the same form as the data was gathered whereas *symbolic* representations extract particular information from the original data source. They also state, “the approach that holds perhaps the most promise for natural and effective awareness support is the situated-literal approach. Here, awareness information is integrated into the workspace’s existing representation, and

is shown in the same form that it was produced by another person. This approach is the closest approximation of how awareness information appears in the real world, and it is the only one that allows people to use their existing skills with the mechanisms of feedthrough, consequential communication, and gestural communication.”

The unfortunate consequence of a literal-situated display is that it inherently limits the ability to collocate personal and shared spaces. Maintaining both individual and group SA will be essential for cross domain military operations. For project UNITY, representational differences (symbolic presentation) in views can come from: *a)* individual tasks, *b)* restrictions arising from data classification, *c)* a user’s taste, or *d)* differences in ethnographic norms. Representational differences can make gestural communication and diectic reference (two major activities in collaboration [10]) difficult to manage. While there can be radical differences in visual representation, a continuous mathematical function that allows for mapping between one domain and another can be created (as depicted in the fisheye lens from Gutwin).

There are pros and cons between pixel-scraping systems and the UNITY-styled implementation as shown below in tables 2 and 3.

Table 2: Pros and cons of pixel-scraping with respect to cross domain collaboration

Cons to pixel-scraping systems	Pros to pixel-scraping systems
<ul style="list-style-type: none"> • Effectiveness is bounded by network bandwidth and latency • Requires a high amount of network bandwidth • Master/slave configuration, no peer-to-peer equality • Cannot be used across trusted cross domain Guards • What I see is limited to what I can share, or what is shared with me • Everyone has to remember everything they cannot say due to classification • Only one system can contribute and multiple peers see identical content • Markups are not embedded into the content and are merely displayed on top 	<ul style="list-style-type: none"> • Add-on software can be used and does not require modification of software that has no collaboration concepts • What I see is what I want you to see and only what I want you to see • Often include capabilities to share documents and provide text chat, and sometimes include voice chat • Source data is not sent in such a way that it can be easily used outside of the sharing session • Often support marking up the scene

Table 3: Pros and Cons of UNITY design

Cons to UNITY design	Pros to UNITY design
<ul style="list-style-type: none"> • Issues can arise when visualization settings or displays are not identical • Naively discussing mixed classification content in your view(s) may disclose non-shared or shareable content • Shared data needs to be resident on both sides, accessible to both sides, or must be distributable across Guards • Referential pointing may fail to create a common item for conversation (items under the mouse may not be the same item, see section 6.1.2) 	<ul style="list-style-type: none"> • Reduced used network bandwidth and latency, due to only sending scene camera information • Complete peers-to-peers system with equal capabilities • Data as presented maintains its form (e.g., a double is a double, and a string is a string) • It's easy to share mouse pointer information, camera location, and row selection in tables • Appropriately marked information can easily be shared to any number of participants • Non-marked data is never shared • Referential pointing can include non-pixel based items, i.e. rows in a table, or items in the view. • Allows for markups to be embedded back into content, however this is not a feature of the current system • Maintains personal SA while still sharing and collaborating with others • Personal preferences which enhance performance can be used to override shared settings • Displaying non-shared data different in-situ with shared data, releases the user from having to remember what cannot be discussed • The amount of data and display real estate is set by the consumer, not the producer

5.2 Use Cases

There are many published Air Force documents that cite the need for collaboration [19][20][9]. The UNITY project therefore had to choose one domain that stressed technical challenges, offered alternative concepts of operations (CONOPS), and was of interest to our partners. All these objectives can be found in the vision of a Combined Space Operations Center (CSpOC), which poses the most challenging problem to address since the domain includes: *a*) an ever increasing number of commercial satellite operators, *b*) national assets from organizations such as the National Aeronautics and Space Administration (NASA) and the European Space Agency (ESA), *c*) US DoD satellites, and *d*) military satellites from other countries. The following is a quote from Ms. Madelyn R. Creedon, Assistant Secretary of Defense Global Strategic Affairs regarding the need for combined space operations.

Led by General Kehler at STRATCOM, the Department is working to transition today's Joint Space Operations Center into a Combined Space Operations Center (CSpOC). A CSpOC will leverage allied space capabilities to augment our own and increase resilience. It will support our ability to operate in coalition operations as we do in other domains and bolster collective defense and deterrence of attack against collective space assets.

Combined space operations require increased sharing of space situational awareness and operational information. Earlier this year, the Secretary of Defense transferred to the Commander of USSTRATCOM his authority to enter into space situational awareness (SSA) data sharing agreements with foreign governments. This compliments USSTRATCOM's existing authority to negotiate SSA sharing agreements with commercial satellite operators. With the extension of this authority to foreign governments, the US will be able to better assist partners with current space operations and lay the groundwork for future cooperative projects.

The increasingly challenging space environment means that an unprecedented level of information sharing is needed among space actors, to promote safe and responsible operations in space and reduce the likelihood of mishaps, misperceptions, and mistrust.[5]

While the variety of stakeholders is not unique to space operations, the type and number of options available in space are greatly restricted. In emergency situations, one could shut off the Internet or close the airspace and force aircraft to land; there is no allegory to that in space. Given the large number of options for experimentation that the CSpOC affords, it was critical to identify a subset of scenarios in order to constrain the possible solution space, and four are presented below.

5.2.1 Master/Slave Use Case (Pixel precise yet not pixel shared)

In this configuration a single producer is in charge of the loading, distributing, and configuring of the visualization. This mode is not strictly enforced in the ASC-UDOP application, but is able to be tested and provided by convention. While this mode seems like an unnecessary substitute to pixel-scraping, recall that cross domain collaboration is not possible at machine speeds with pixel-scraping systems. In this mode, all the data brought into the system is considered shareable and transmitted across the guard. This is also a particularly attractive solution when the data is not frequently updated and where having each client construct the graphics locally provides for a significant reduction in network bandwidth.

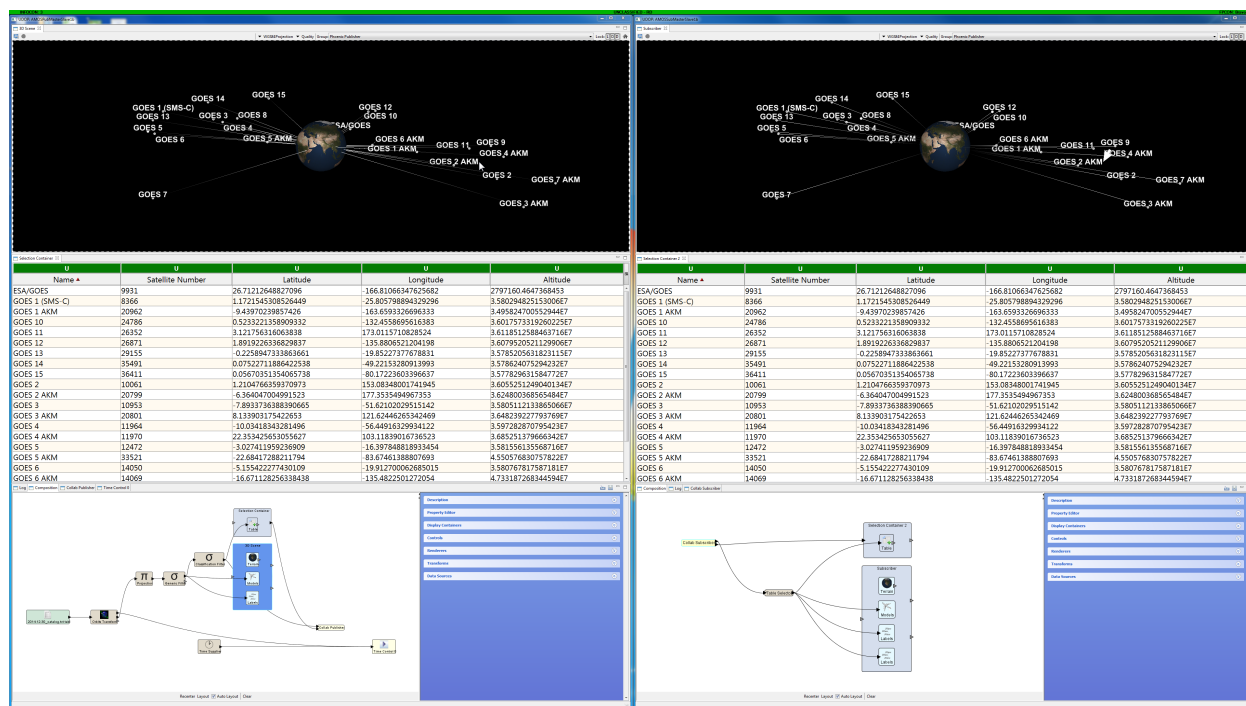


Figure 4: Displaying all the satellites whose names start with GOES and sharing (from the master on left) content with a single client (on right).

Notice that the composition configuration (the lower left window in the master and slave) is different between the master and the slave visualization. In this case, the master (left) loaded the satellites Two Line Element (TLE) information, propagated the orbits, added artificial classification, performed a generic filter satellites that did not start with the word 'GOES', then filters them based on classification, sends that data to the 2D table and to potential subscribers via the Collab Publisher component.

The client (right) composition window is a Collab Subscriber and selects the published table of interest (in this case

only one is being published) and generates an identical visualization. You should also note that the mouse cursor present on the master is being represented by a mirrored pink version on the client.

5.2.2 Masters/Slaves Use Case

The ASC-UDOP capabilities allow for a unique mode where there can be multiple masters and multiple slaves. In the traditional pixel-scraping system (referred here as a Master/Slave configuration) there exists a single source of information. In the Masters/Slaves mode, each user can bring their own data into the system and control how it is presented. This still provides for symmetrical views (situated-literal) but is a conglomeration of multiple sources of data (Masters) which can be used to greatly enhance shared situation awareness. This is a significant improvement over the capabilities of pixel-scraping systems independent of the fact that pixel-scraping cannot be interrogated by Guards and is not a viable option currently.

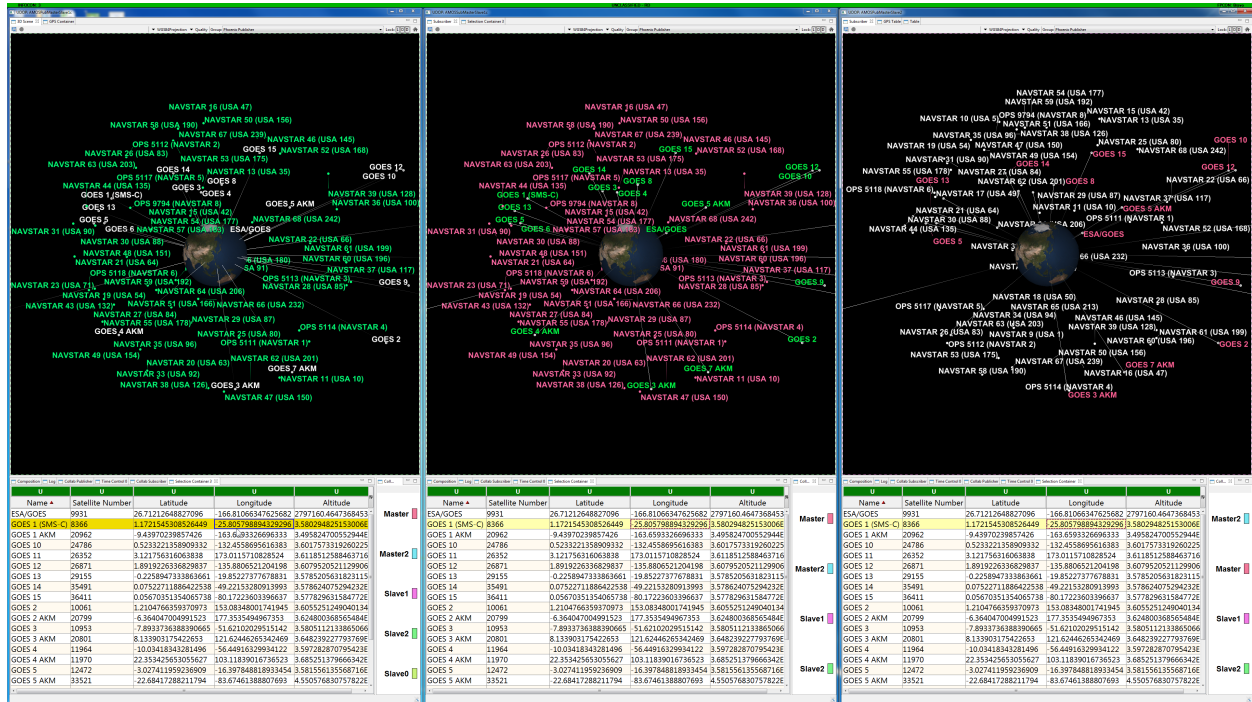


Figure 5: Two Masters (left and right) with one slave (middle). In this case the Master (left) is contributing the GOES satellite constellation and the Master (right) is contributing the GPS constellation information. Each client is configured to display the information they are providing as white, data from Master (left) is displayed as fuchsia, and data from Master 2 (right) is displayed as green.

5.2.3 Peer to Peer Use Case

In this configuration each user can have a superset of the shared data and may have different visual configurations for the same item. Just how *Relaxed* can a RWYSIWIS visualization be before all communication breaks down and shared awareness is lost? The UNITY project doesn't answer that question, but does produce a platform for that experimentation, see Figures 6 and 7.

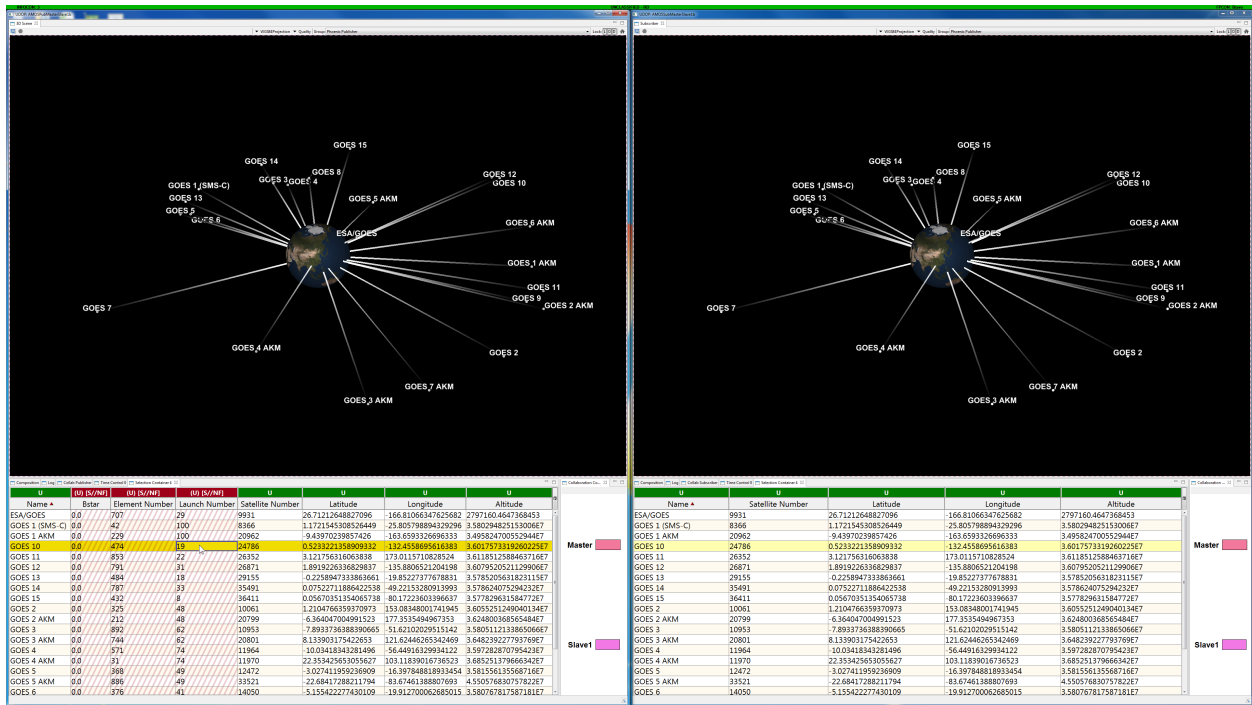


Figure 6: The synthetically classified data (Bstar, Element number, and Launch Number) are displayed in the master but not being transmitted to the slave. Even though there are a different number of columns available to the master, the row containing the mouse pointer is represented in the client display as a highlighted row in the client display.

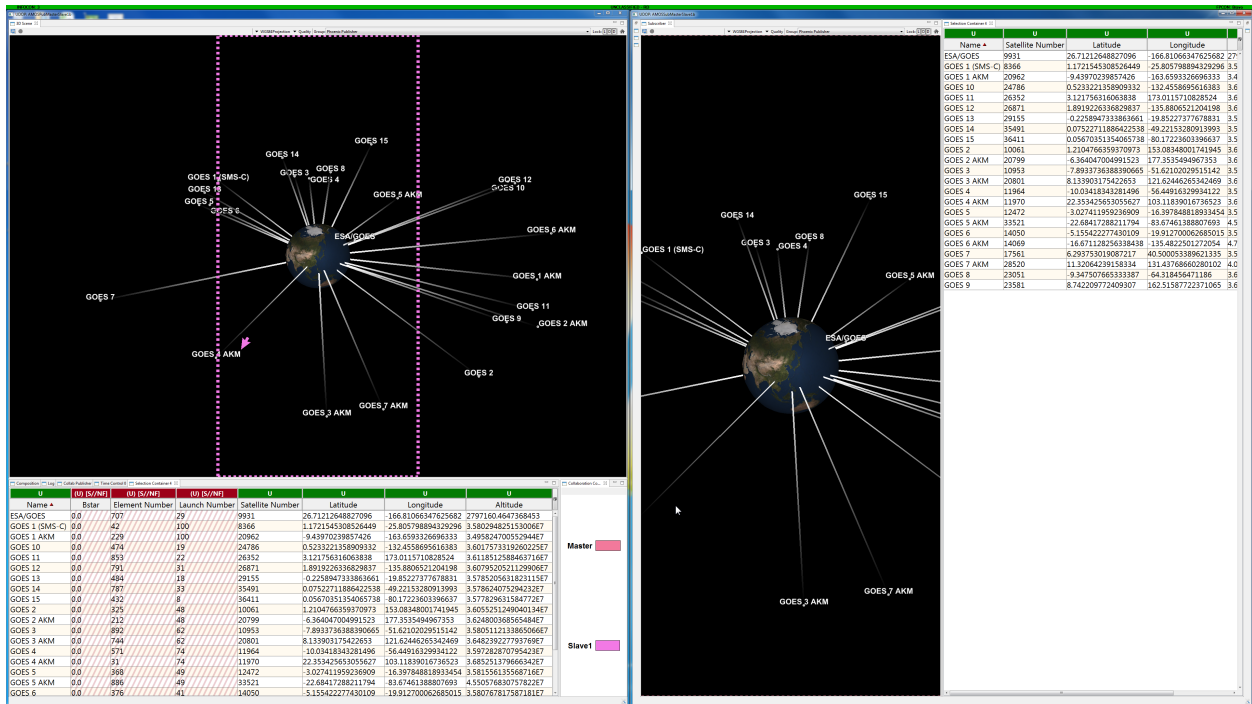


Figure 7: Notice the different window configuration and the dashed lines on the master (left) showing the extents of the client's 3D View. Notice that the mouse cursor is visible on the client's (right) 3D View. Additionally, notice that the label for GOES 2 AKM is absent from the client; this is discussed in the Text Decluttering section below.

5.2.4 Peers to Peers Use Case

This is very similar to the Peer to Peer use case, but has an interesting distinction. Every operator may be sharing content with a set of participants and not the other(s) while still having a locally combined view of all the data being shared with them. Making the source content of any one visualization a combination of asymmetrically shared content among all simultaneous collaborators. Of course each user retains the ability to have different visual configurations for the same item, possibly resulting in a maddening array of options (see Figure 8).

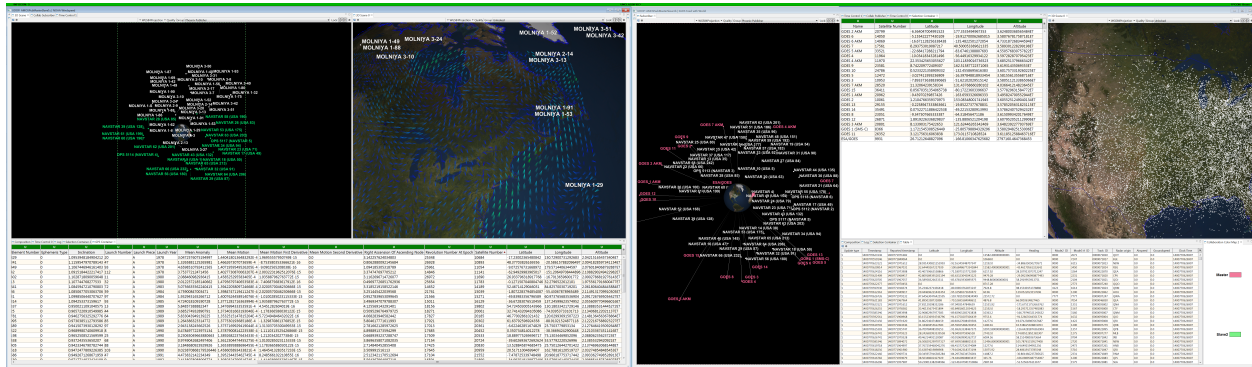


Figure 8: Peer1 (left) is displaying a synchronized 3D View with the data being provided from Peer2 (right) with additional information and has a completely non-shared 3D view of the weather information with a complete list of satellites in the Tabular View below. Peer2 is displaying a superset of satellite information, some of which was shared by Peer1 and then simultaneously viewing the active air picture over the United States with a table of that content.

Table 4: “Elements of workspace awareness relating to the present” from [11] and ASC-UDOP’s treatment

Category	Element	Specific questions	ASC-UDOP treatment
Who	Presence	Is anyone in the workspace?	Currently displayed in the connection dialog and presented as an optionally added user list.
	Identity	Who is participating? Who is that?	Currently embodied as a unique color for each cursor in the 3D views and highlighted cell in the Tabular view.
	Authorship	Who is doing that?	Currently not represented specifically in the Graphical User Interface (GUI), but the user could visually represent data from each provider differently based upon the need.
What	Action	What are they doing?	This is supported by the location of the other users’ mouse pointers in the 3D View and the Tabular view.
	Intention	What goal is the action part of?	This is well supported in the 3D view with the shared cursors, but is very complicated and discussed in detail in the Tabular views section.
	Artifact	What object are they working on?	In traditional WYSIWIS systems this refers to a collaborant manipulating something in the shared view, such as a value in a cell or an icon on a map. This is not treated in ASC-UDOP due to the fact that content is coming from datasources, and not the user directly.
Where	Location	Where are they working?	This is supported by view rectangles in the 3D view, but other portions of the display have no natural opportunities for the same capability.
	Gaze	Where are they looking?	This is done through the assumption that the mouse cursor location or cell highlight represents the gaze location.
	View	Where can they see?	This is only supported in the 3D view, not the Tabular view.
	Reach	Where can they reach?	This is supported through view synchronization in the 3D View.

6. UNITY COLLABORATIVE VIEWS

While there are many visual components in the ASC-UDOP, collaboration features have only been added to the 3D View and 2D Table. Other components (graphs, charts, and the composition window) could be enabled and may offer unique challenges and opportunities but are beyond the scope of this effort. Table 4 describes some of the elements of workspace awareness and how they are treated in the current ASC-UDOP application.

6.1 3D View Collaboration

The components of 3D view collaboration consist of synchronization of the following: the view location, the data, and visualization settings.

A shared 3D View may have the same or different:

- display width and height in pixels,
- display area aspect ratio,
- viewing vector,
- viewing orientation,
- frustum (described below),
- formatting for the contents in the view, and
- mouse location.

For shared geospatial situation awareness, a common technique is to present a copy of the same view to all participants using pixel-scraping software. However, recall that, in the UNITY experiment, the exact pixels of the image are not shared. Therefore, to create a similar capability to pixel-scraping, the camera location information is shared to participants so the client(s) can re-create the same or similar (e.g., different canvas size or alternative scene contents) scene view(s) locally. By not sharing the exact pixels, UNITY allows multiple participants to display different information.

While there could also be differences in physical device sizes that could alter each user's ability to perceive shared content, those issues are out of scope for this effort but do pose interesting future opportunities.

In 3D graphics, there are six planes that define the volume of space that is projected onto the screen. These planes as a whole are called a frustum. Since UNITY allows for asymmetries between each users' composition of views, the corresponding frustums can also have the same or different: aspect ratio, near and far clipping plane distances from the center location, and center location itself. (see Figure 9).

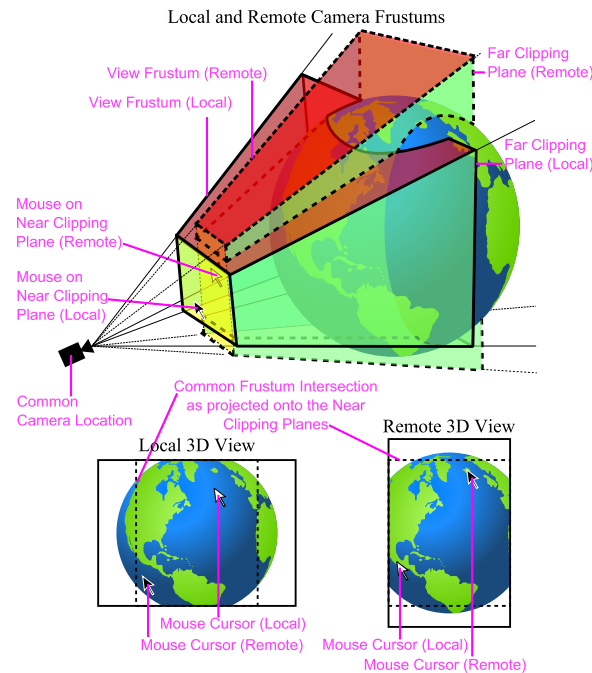


Figure 9: Local and Remote 3D Visualization, the subset of view frustum and canvas size mis-mappings that are supported by UNITY

There are actually many more mis-alignments of frustum configurations that could be supported, and some are presented in Figure 10. This figure only displays the mis-mappings that would result in a 2D misalignment of content. Of course a complete 3D mis-alignment of frustums is possible and attempting to show each participant's unique 3D viewing volume in an understandable fashion is well-beyond the scope of this project.

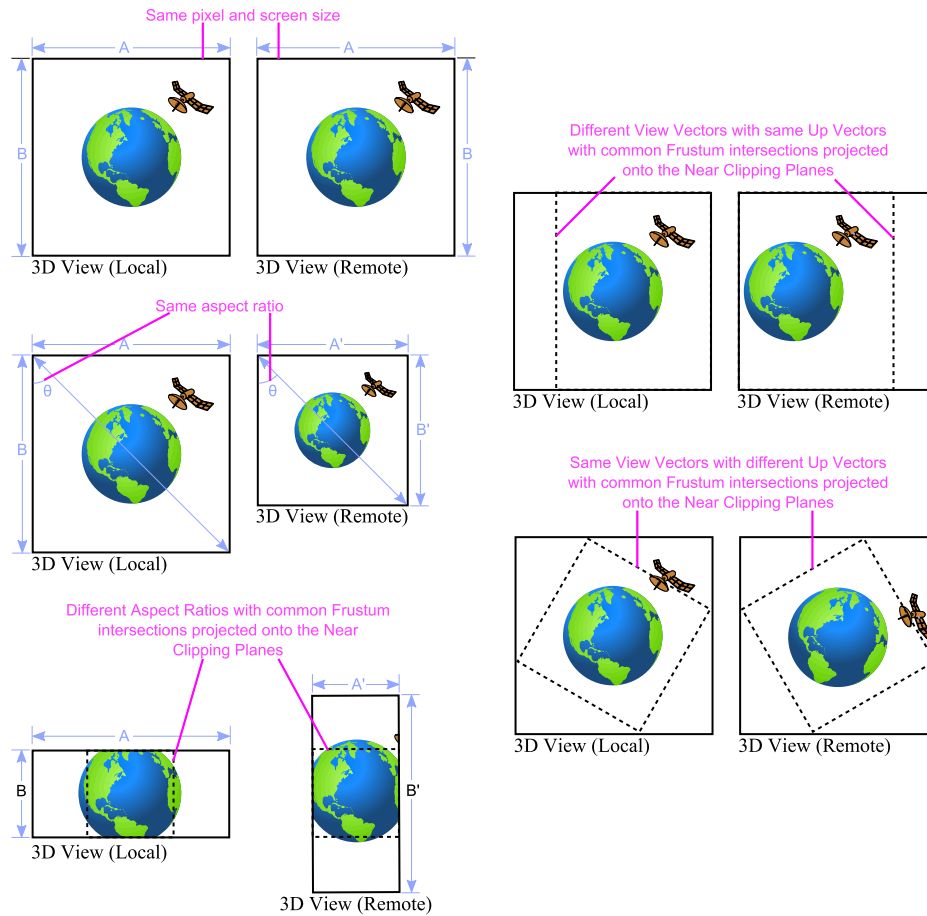


Figure 10: Local and Remote 3D Visualization, a subset of view frustum and canvas size mis-mappings

6.1.1 View Synchronization

View synchronization support existed in the ASC-UDOP application prior to the UNITY effort. If the user had multiple views open, they could synchronize the camera angles among those views. The ASC-UDOP (prior to UNITY) presented the user with the graphical controls (see left in Figure 11). Any views in the positive numbered groups (1, 2, or 3) were kept in sync with respect to the camera locking parameters of Location (L), Orientation (O), and Distance (D). Group 0 was reserved for views that were not synchronized.

While all the options (view groups and camera synchronization parameters) are still relevant for local and remote collaboration, their names needed to become more explicit to differentiate between local and remote collaboration. This also provides the opportunity to restrict certain options, remote mouse collaboration between views is only supported when the location, orientation, and distance parameters are completely synchronized.




Figure 11: Comparison between old (left) and new (right) view synchronization controls

6.1.2 Text Rendering in 3D Applications and Text Decluttering

An integral portion of many visualizations is the effective use of labels. The ASC-UDOP application extended the particle-based labeling algorithm by Luboschik et al. [14] by incorporating frame to frame coherency similar to [6]. Some important traits for integration into the ASC-UDOP application were runtime performance, deterministic label placement, and the ability to expose numerous tunable parameters such as label priority. These three traits are critical for users to understand the dynamics of the display system while still being able to modify them to suit numerous operational domains.

The user has fine control over many label parameters: label anchor location, maximum and minimum label offset from the anchor, the font color, the font style, the font family, the priority for occlusion determination, and many others. However, even with all the options, there are two problems that still emerge even if all the settings for the labeling feature are the same between 3D views, drastically affecting label location and visibility. The first problem comes from the desire to have text readable by the user. In many systems, as is the case with the ASC-UDOP, text is rendered in screen space. This means that text does not change size as the viewer's distance from the labeled objects change. While this is normally a highly desirable trait, collaborative 3D views as in Figure 10 are not required to be of identical pixel size or aspect ratio and, along with the dynamic text decluttering algorithm, are the root causes for many problems.

Problem 1, since the label locations are deterministic and processed in priority order, a partially obscured label may be forced on screen, resulting in a very different final position for every label with a lower priority. Problem 2, the second problem can emerge even if all the visual settings for the 3D views and the label parameters are identical. This stems from the fact that during collaboration when a subset of data is chosen to be transmitted to a partner this will cause a different number of labels to be displayed. And therefore, the label's placement can become radically different as it avoids overlapping text.

The first thought may be to globally set the label locations from one 3D view and have the others display the labels in the screen locations as defined by the master, even if the master has more labels. This unfortunately could lead to a very awkward view for the client (some labels may be drawn very far away when they could have been drawn closer) and exposes a potential security problem. Since the algorithm is deterministic, any non-consistent label placement may indicate that certain information is being withheld. This may be interpreted as a lack of trust even when the data is being withheld just for clarity and to focus on a particular subset of all the available data.

There is a relatively simple solution when 3D views have identical sizes and label parameters. To solve it, the label placement priority can be specified as a column in the data table (not the visual Tabular View, but rather the internal tabular data structure) and be used to create consistent views among participants, even if one user has additional labels. This has the side effect of the additional labels are always subordinate to the shared labels and may not be drawn since they are occluded (see Figure 12).

6.2 Tabular View Collaboration

While the 3D view is challenging with respect to all the visual mis-mapping problems that can occur, it does not expose every possible asymmetric variant possible; the ASC-UDOP tabular view (see Figure 13) presents even more problems.

While the values of the data depicted in the table aren't directly manipulable, numerous view options are configurable by the users. The list includes: column visibility, column order, sorting order for values within each column (including hierarchical sorting), table font family, table font size, table font style, highlighted row color, selected row color, compound highlighted and selected row color, and for appropriate schema formats: degrees/minutes/seconds or heading arrows. While that list may seem exhaustive, it does not address the issue of additional transformations being added to the data flow diagram that modify the values within a collaborant's view.

The important elements that affect the table are the following; a shared table may have a different: display width or height in pixels, number of rows, number of columns, formatting for the elements in a row or column, sorting criteria,



Figure 12: Comparison of text decluttering priority with different label sets. The GOES satellite labels in priority order based on their satellite number (left). All the satellites with the GOES satellites priority based on their satellite number and all other satellites priority based below the GOES on depth from the viewer (middle). All the satellites without assigning priority (right), causing the shared labels to be occluded.

and mouse location. It's also important to remember the initial requirement of shared pointers needs to be taken into account for each of these options. Another factor is what portions of the table are visible when scrolling can occur independently, see Figure 13.

In order to address the asymmetry on releasability, each column now has a 'Classification//Release Marking' field added in accordance to published doctrine and policy[17]. While the vivid red background behind the marking label for classified SECRET data is not strictly law, it is considered a best practice and continued here in this work. The marking created an additional constraint on the user experience since it always has to be visible. Portion marking guidance states that, 'every portion in every classified document shall be marked to show the highest level of classification that it contains... the criterion [for marking subportions is to] avoid over-classification of any of the information or eliminate doubt about the information's classification level[18]. The overall classification marking in the application window shows the maximum classification of data in the system, and for the UNITY experiments all the markings are for test only, and do not represent real world systems or classifications; all the content is UNCLASSIFIED.

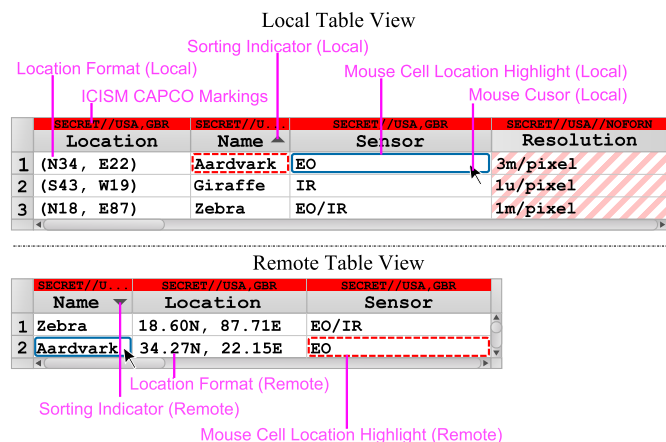


Figure 13: Local and Remote Tabular Visualization, a subset of visual mis-mappings

7. CONCLUSION

The UNITY program successfully created a flexible platform for a collaborative Relaxed What You See Is What I See environment in compliance with DoD policy. It is clear that collaboration will only increase among disparate users as budgets shrink, expertise becomes more distributed, and team construction includes unconventional (and temporary) partners. Addressing the security implications by reducing functionality will only hurt the DoD and the UNITY platform shows that Situation Awareness can be maintained or increased even in an environment of unbalanced information availability. The unique capabilities of the UNITY concept resulted in a powerful set of capabilities, a Federated Collaborative User-Defined Operating Picture (FC-UDOP).

8. FUTURE WORK

The UNITY International Agreement continues until 2017, and the Air Force Research Laboratory is working on relevant space operation concepts in conjunction with the United Kingdom Space Operations Coordinations Center. While the initial set of experiments is more about testing the technology and long haul collaboration, the human aspect of the work has numerous opportunities to inform the capabilities that should and should not be exposed to users. One interesting study is whether long haul communication with and without speech can maintain the normal social contract of when one is speaking that others listen. With the latency that may be inherent in the system and a lack of active locking, the users may experience a similar problem that is present when people try to voice chat with a delay and talk over one another. In the worst case with multiple participants this could devolve into chaos and force the users into a master/slave configuration.

Another interesting avenue to research is in ensuring that the values being displayed and discussed are either mathematically are exactly the same. For instance, the US operator may want to see units in feet, where the UK operator may prefer to see them in meters. To address this issue, popups, or additional columns could be used to augment the information and in some way tie them together visually so that highlighting your new visual representation still results in highlighting the correct alternative local and remote view components.

It is critical that a properly formed human experiment be performed to measure any effect. It may turn out that while this design supports real-time uninterrupted collaboration, the complications that it creates prove to be too numerous. In that case, the UNITY technology could revert to a Master/Slave configuration that still allows for a single system connected to multiple clients across security domains.

The UNITY concept is predicated on accepting that clients will be receiving data in a form in which standard network sniffing software could grab and potentially use the data for longer than desired by the producer. Another interesting alternative to this would be to add a remote rendering capability after the Guard and only ship the pixels as is done in standard pixel-scraping software. This would allow the Master to retain the ability to have content displayed within their view that is different with what is being shared and also prevent the client(s) to easily scrape the data for later use.

Nearly all visual components within the ASC-UDOP are mappable to input parameters coming from data sources and are scriptable by the user using Groovy or Python. This means that the operators have the freedom, and responsibility, to make the visualization as they see fit to meet the current need. This fits well with one of our mantras, 'The science of control shall not inhibit the art of command.' While that flexibility is paramount to adjusting in real-time as a situation unfolds, it almost encourages visual mappings to be different among collaborators. This asymmetry between views may produce more barriers than it provides and determining the Tactics, Techniques, and Procedures (TTPs) for such a flexible system need to be determined.

For instance, the system may need to employ semantic coloring versus exact coloring. The color of an icon for military displays is extremely important and, for the US, is detailed by the MilStd-2525 symbology specifications[7]. This specification maps friendly forces to blue rectangles, enemy forces to red diamonds, neutral forces to green squares, and unknown forces to yellow clover leaves. This makes sharing a visualization a matter of either sharing the semantics of being friendly or the color of the icon being blue. This could make collaboration between these participants when augmented by secure 2 way voice communication even more difficult. When one participant speaks about a blue icon on the his or her screen, the actual color on the other collaborant's screen could be vastly different.

ACKNOWLEDGMENTS

The authors wish to thank the UNITY Program Manager (PM) Mr. Tony DiLego, the Phoenix' PM Mr. Timothy Blocher, the CACI development team who developed the preponderance of the software, the UK Ministry of Defence, the US Department of Defense, and the International Affairs Office of the Secretary of the Air Force for the opportunity to work in these challenging environments and for the ability to publish the work to the larger research community.

References

- [1] AFRL. Air, space, cyber user defined operational picture (asc-udop). <http://www.wpafb.af.mil/shared/media/document/AFD-110516-041.pdf>.
- [2] AFRL. Phoenix prime, universal information management services for all domains. http://www.dodsbir.net/sitis/view_pdf.asp?id=AF141_036_Ref_3_AFRL_Phoenix_Prime_Fact_Sheet.pdf.
- [3] AFRL. User needs drive air, space, cyber - user defined operational picture. <http://www.wpafb.af.mil/news/story.asp?id=123296337>.
- [4] R. Bosch, C. Stolte, D. Tang, J. Gerth, M. Rosenblum, and P. Hanrahan. Rivet: A flexible environment for computer systems visualization. *ACM SIGGRAPH Computer Graphics*, 34(1):68–73, 2000.
- [5] S. A. S. Committee. Statement of ms. madelyn r. creedon assistant secretary of defense global strategic affairs before the senate committee on armed services subcommittee on strategic forces, Mar. 2012. http://www.defense.gov/home/features/2011/0111_nsss/docs/ASD%20Madelyn%20Creedon%20Written%20SASC%20Space%20Testimony%20March%202012.pdf.
- [6] H. Cords, M. Luboschik, and H. Schumann. Floating labels: Improving dynamics of interactive labeling approaches. In *Proceedings of MCCSIS (IADIS Multi Conference on Computer Science and Information Systems)*, pages 235–238, 2009.
- [7] Department of Defense. *Common Warfighting Symbolology*, revision c edition, Nov. 2008. http://www.assistdocs.com/search/document_details.cfm?ident_number=114934.
- [8] Exelisinc.com. Exelis - isse (information support server environment), 2015. <http://www.exelisinc.com/solutions/ISSE/Pages/default.aspx>.
- [9] U. S. A. Force. *Space Operations, Air Force Doctrine Document 3-14*, June 2012. http://static.e-publishing.af.mil/production/1/af_cv/publication/afdd3-14/afdd3-14.pdf.
- [10] C. Gutwin and S. Greenberg. Design for individuals, design for groups: tradeoffs between power and workspace awareness. In *Proceedings of the 1998 ACM conference on Computer supported cooperative work*, pages 207–216. ACM, 1998. <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.39.3919&rep=rep1&type=pdf>.
- [11] C. Gutwin and S. Greenberg. A descriptive framework of workspace awareness for real-time groupware. *Computer Supported Cooperative Work (CSCW)*, 11(3-4):411–446, 2002. <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.330.8099&rep=rep1&type=pdf>.
- [12] M. G. J. A. Hawkins. *DEFENSE INFORMATION SYSTEM NETWORK (DISN): POLICY, RESPONSIBILITIES AND PROCESSES*, July 2003. http://jitc.fhu.disa.mil/jitc_dri/pdfs/6211_02b.pdf.
- [13] J. Li, S. Greenberg, E. Sharlin, and J. Jorge. Interactive two-sided transparent displays: Designing for collaboration. 2014. <http://grouplab.cpsc.ucalgary.ca/grouplab/uploads/Publications/Publications/2014-TransparentDisplay-Report2014-1053-03.pdf>.
- [14] M. Luboschik, H. Schumann, and H. Cords. Particle-based labeling: Fast point-feature labeling without obscuring other visual features. *Visualization and Computer Graphics, IEEE Transactions on*, 14(6):1237–1244, 2008.

- [15] A. McVay, D. Krisher, and P. Fisher. Jview visualization for virtual airspace modeling and simulation. Technical report, Air Force Research Laboratory, 2009. <http://www.dtic.mil/cgi-bin/GetTRDoc?AD=ADA496743>.
- [16] A. McVay, D. Krisher, and P. Fisher. Jview visualization for next generation air transportation system. Technical report, Air Force Research Laboratory, Jan. 2011. <http://www.dtic.mil/cgi-bin/GetTRDoc?AD=ADA538715>.
- [17] P. B. H. Obama. *Executive Order (EO) 13526, Classified National Security Information*, Dec. 2009. <http://www.ise.gov/sites/default/files/EO13526.pdf>.
- [18] D. of Defense. *DoD Information Security Program: Marking of Classified Information*, Mar. 2013. http://www.dtic.mil/whs/directives/corres/pdf/520001_vol2.pdf.
- [19] D. of Defense. *Command and Control of Joint Air Operations*, Feb. 2014. http://www.dtic.mil/doctrine/new_pubs/jp3_30.pdf.
- [20] D. of Defense. *Basic Doctrine*, Feb. 2015. <https://doctrine.af.mil/download.jsp?filename=Volume-1-Basic-Doctrine.pdf>.
- [21] G. E. Seymour and M. B. Cowen. A review of team collaboration tools used in the military and government. Technical report, DTIC Document, 2007.